

08/07/00
JC891 U.S. PTO

08-09-80

A+

PTO/SB/05 (2/98) (modified)
Approved for use through 9/30/2000, OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

NEW UTILITY PATENT APPLICATION TRANSMITTAL <i>(only for new nonprovisional applications under 37 CFR 1.53(b))</i>	Attorney Docket Number	4466 US
	First Named Inventor	Pawan Goyal
	Total Pages in this Submission	41
	Express Mail Label No.	EL482474496US

JC896 U.S. PTO
09/633575

08/07/00

APPLICATION ELEMENTS	ACCOMPANYING APPLICATION PARTS
1. <input checked="" type="checkbox"/> Fee Transmittal Form (in duplicate) <input checked="" type="checkbox"/> Checks Enclosed (\$345.00 / Filing Fee & \$123.00 / Extra Claims) 2. <input checked="" type="checkbox"/> Specification <i>(preferred arrangement set forth below)</i> <input type="checkbox"/> Descriptive Title of the Invention <input type="checkbox"/> Cross Reference(s) to Related Case(s) <input type="checkbox"/> Statement Regarding Fed sponsored R & D <input type="checkbox"/> Background of the Invention <input type="checkbox"/> Brief Summary of the Invention <input type="checkbox"/> Brief Description of the Drawing(s) <input type="checkbox"/> Detailed Description <input type="checkbox"/> Claim or Claims <input type="checkbox"/> Abstract of the Disclosure 3. <input checked="" type="checkbox"/> Drawing(s) (when necessary per 35 USC 113) 4. Oath or Declaration a. <input checked="" type="checkbox"/> New Declaration <input checked="" type="checkbox"/> Executed b. <input type="checkbox"/> Copy from a prior application (37 CFR 1.63(d)) <i>(for continuation/divisional with Box 17 completed)</i> i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b). 5. <input type="checkbox"/> Incorporation by Reference (useable if Box 4b is checked). The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.	6. <input type="checkbox"/> Assignment & Assignment Recordation Cover Sheet 7. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>(if foreign priority is claimed)</i> 8. <input type="checkbox"/> Information Disclosure Statement & PTO-1449 <input type="checkbox"/> Copies of IDS Citation(s) 9. <input type="checkbox"/> Preliminary Amendment 10. Small Entity Statement <input checked="" type="checkbox"/> New Statement enclosed <input type="checkbox"/> Statement filed in prior application. Status still proper and desired 11. <input checked="" type="checkbox"/> Return Postcard 12. <input type="checkbox"/> _____ 13. <input type="checkbox"/> _____ 14. <input type="checkbox"/> _____ 15. <input type="checkbox"/> _____ <div style="text-align: center;"> ADDRESS TO: Commissioner for Patents Box Patent Application Washington, D.C. 20231 </div>
17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment: <input type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No: ____/____ Prior application information: Examiner: _____ Group/Art Unit: _____	

18. CORRESPONDENCE ADDRESS					
NAME	Renée M. DuBord, Esq. Fenwick & West LLP				
ADDRESS	Two Palo Alto Square				
CITY	Palo Alto	STATE	CA	ZIP CODE	94306
COUNTRY	U.S.A.	TELEPHONE	(650) 858-7880	FAX	(650) 494-1417
Name (Print/Type)	Renée M. DuBord, Esq.		Registration No. (Attorney/Agent)		42,500
Signature	<i>Renée DuBord</i>		Date	8/7/00	

EL482474496US

0002/PTO(modified) U.S. Department of Commerce
Rev. 10/95 Patent and Trademark Office

FEE TRANSMITTAL

TOTAL AMOUNT OF PAYMENT

Subtotal (1) + Subtotal (2) + Subtotal (3) = **\$468.00**

Complete if Known

Application Number	Not Yet Assigned
Filing Date	August 7, 2000
First Named Inventor	Pawan Goyal
Group Art Unit	Not Yet Assigned
Examiner Name	Not Yet Assigned
Attorney Docket Number	4466 US

METHOD OF PAYMENT

1. The Commissioner is hereby authorized to:

- ☐ Charge the indicated fee (Small Entity Filing Fee) to the below mentioned deposit account.
- ☒ Charge any additional fee required under 37 CFR 1.16 - 1.21 or credit any over payments to the below mentioned deposit account. †
- ☐ Charge the Issue Fee set in 37 CFR 1.18 at the Mailing of the Notice of Allowance, 37 CFR 1.311(b) to the below mentioned deposit account.

Deposit Account Number: 19-2555

Deposit Account Name: FENWICK & WEST LLP

A Duplicate Copy of this authorization is attached

2. ☒ Payment Enclosed:

☒ Checks ☐ Other

FEE CALCULATION (fees effective 11/12/98)

1. FILING FEE

Large Entity Fee Code/Fee	Small Entity Fee Code/Fee	Fee Description	Fee Due
101/\$690	201/\$345	Utility Filing	345
102/\$310	206/\$155	Design Filing	
108/\$690	208/\$345	Reissue	
114/\$150	214/\$75	Provisional Filing	
SUBTOTAL (1)			(\$345.00)

2. CLAIMS

Large Entity Fee Code/Fee	Small Entity Fee Code/Fee	Fee Description
103/\$18	203/\$9	Claims in excess of 20
102/\$78	202/\$39	Independent claims in excess of 3
104/\$260	204/\$130	Multiple dependent claim
109/\$78	209/\$39	Reissue independent claims over original patent
110/\$18	210/\$9	Reissue claims in excess of 20 and over original patent

3. ADDITIONAL FEES

Large Entity Fee Code/Fee	Small Entity Fee Code/Fee	Fee Description	Fee Due
105/\$130	205/\$65	Surcharge - late filing fee or oath	
127/\$50	227/\$25	Surcharge-late provisional filing fee or cover sheet	
147/\$2,520	147/\$2,520	For filing a request for reexamination	
115/\$110	215/\$55	Extension for response within first month†	
116/\$380	216/\$190	Extension for response within second month†	
117/\$870	217/\$435	Extension for response within third month†	
118/\$1,360	218/\$680	Extension for response within fourth month†	
128/\$1,850	228/\$925	Extension for response within fifth month†	
119/\$300	219/\$150	Notice of Appeal	
141/\$1,210	241/\$605	Petition to revive unintentionally abandoned application	
142/\$1,210	242/\$605	Utility Issue Fee (Or Reissue)	
143/\$430	243/\$215	Design Issue Fee	
122/\$130	122/\$130	Petitions to the Commissioner	
123/\$50	123/\$50	Petitions related to provisional applications	
126/\$240	126/\$240	Submission of Information Disclosure Statement	
581/\$40	581/\$40	Recording each patent assignment per property (times number of properties)	
146/\$690	246/\$345	Filing a submission after final rejection (37 CFR 1.129(a))	
149/\$690	249/\$345	For each additional invention to be examined (37 CFR 1.129(b))	
Other fee (specify):			
Other fee (specify):			
SUBTOTAL (3)			(\$0.00)

(Col. 1)		(Col. 2)		(Col. 3)		Fee Due	
For	No. of Existing Claims	Highest No. Previously Paid For		Extra**	Fee		
TOTAL	25	20 or 0	=	5	9	=	45
INDEP	5	3 or 0	=	2	39	=	78
[] First presentation of multiple dependent claim							

* Subtract the greater number of Col. 2

** If the difference between Col. 1 and Col. 2 is less than zero, then enter "0" in Col. 3

SUBTOTAL (2) \$123.00

SUBMITTED BY		Complete (if applicable)	
Typed or Printed Name	Reneé M. DuBord	Reg. Number	42,500
Signature	<i>Reneé DuBord</i>	Date	8/7/00



VERIFIED STATEMENT CLAIMING SMALL ENTITY STATUS (37 CFR 1.9(f) & 1.27(c))--SMALL BUSINESS CONCERN	Docket Number (Optional): 4466 US
--	--------------------------------------

Applicant or Patentee: Pawan Goyal and Srinivasan Keshav
Application or Patent No.: not yet known
Filing Date or Issue Date: not yet known

Title: Fairly Partitioning Resources While Limiting The Maximum Fair Share

I hereby declare that I am

☐ the owner of the small business concern identified below:

☒ an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF SMALL BUSINESS CONCERN Ensim Corporation

ADDRESS OF SMALL BUSINESS CONCERN 1366 Borregas Avenue, Sunnyvale, CA 94089

I hereby declare that the above identified small business concern qualifies as a small business concern as defined in 13 CFR 121.12, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees to the United States Patent and Trademark Office, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention described in:

☒ the specification filed herewith with title as listed above.

☐ the application identified above.

☐ the patent identified above.

If the rights held by the above identified small business concern are not exclusive, each individual, concern or organization having rights in the invention must file separate verified statements averring to their status as small entities, and no rights to the invention are held by any person, other than the inventor, who would not qualify as an independent inventor under 37CFR 1.9(c) if that person made the invention, or by any concern which would not qualify as a small business concern under 37 CFR 1.9(d), or a nonprofit organization under 37 CFR 1.9(e).

Each such person, concern or organization having any rights in the invention is listed below:

☒ No such person, concern, or organization exists.

☐ Each such person, concern or organization is listed below:

Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING Srinivasan Keshav

TITLE OF PERSON IF OTHER THAN OWNER Chief Technical Officer

ADDRESS OF PERSON SIGNING 1366 Borregas Avenue, Sunnyvale, CA 94089

SIGNATURE [Signature] DATE 8/2/00

**FAIRLY PARTITIONING RESOURCES WHILE LIMITING
THE MAXIMUM FAIR SHARE**

INVENTORS

Pawan Goyal

Srinivasan Keshav

Prepared by:

Renee M. DuBord

Reg. No. 42,500

Fenwick & West LLP

Two Palo Alto Square

Palo Alto, CA 94306

Express Mail No.: EL482474496US

FAIRLY PARTITIONING RESOURCES WHILE LIMITING THE MAXIMUM FAIR SHARE

INVENTORS

5 Pawan Goyal and Srinivasan Keshav

BACKGROUND

Field of Invention

The present invention relates generally to resource scheduling, and more particularly, to scheduling a resource fairly while preventing resource users from exceeding a maximum resource allotment.

Background of the Invention

A resource scheduler performs the function of allocating resources. Different resource types may use separate resource schedulers. Within each resource scheduler, each customer or user of the resource is treated as a separate “schedulable entity” with a separate scheduling queue and an individual quality of service guarantee. The scheduler selects requests for service from
15 among the different queues, using a scheduling algorithm to ensure that each queue receives a least the minimum level of service they were guaranteed. Different queues may have different minimum quality of service guarantees, and the resource requests from each queue are weighted by the quality of service guarantee. Weighting increases or decreases a schedulable entity’s
20 relative resource share.

The goal of the resource scheduler is twofold. First, the resource scheduler must try to ensure that each schedulable entity is allocated resources corresponding to at least the minimum quality of service resource level paid for by that schedulable entity. However, if extra resources are available, the resource scheduler must decide how to allocate the additional resources.

5 Typical resource scheduling algorithms and methods are work-conserving. A work-conserving scheduler is idle only when there is no resource request to service. Additional resources will be divided up among the schedulable entities with outstanding requests, in proportion to each schedulable entity's weight. Customer requests are serviced if the resource is available, even if they exceed the schedulable entities' quality of service guarantee.

Thus, a work-conserving resource scheduler may often provide schedulable entities with service beyond the actual maximum quality of service paid for by the schedulable entity. Unfortunately, this behavior tends to create unrealistic expectations. For example, assume customers A and B both are sharing a resource. Customer A pays for 50% of the resources and customer B pays for 25%, and resource sharing is weighted between A and B to reflect these different allocations. In a work-conserving scheduler, customers with unsatisfied requests get resource shares in proportion to their weights. Therefore, if both A and B request resources beyond their paid-for quality of service guarantee, one embodiment of a work-conserving scheduler will allocate two-thirds of the extra 25% of the resources to A, and the remaining one-third to B.

20 However, if a new customer C is added to further share the resource, and C pays for and receives 25% of the resources, the resources actually delivered to A and B will decrease. Although A and B will still receive the 50% and 25%, respectively, of resources that they paid for, A and B may both perceive a decrease in service. In order to avoid setting unrealistic

customer expectations, it is preferable for the resource scheduler to prevent customers from receiving more resources than they have paid for, even if this allows resources to be idle during certain times. Such a preferred resource scheduler is non-work-conserving, and implements both minimum and maximum quality of service guarantees.

5 One existing method for preventing schedulable entities from exceeding their maximum allotted resources is to place a rate controller module into the system before the resource scheduler. The rate controller module receives resource requests from each different schedulable entity and separates them into individual “schedulable entity” queues. Each schedulable entity has a separate queue. Before a request is passed on to the scheduler, the rate controller module checks to determine if granting the request will exceed the schedulable entity’s maximum resource allotment. If so, the rate controller module sets a timer to expire when the request may be granted without exceeding the schedulable entity’s maximum allotment. Upon timer expiry, the rate controller module allows the request to be passed on to the scheduler, where it will be scheduled for service.

20 The implementation of the rate controller module requires a separate queue and timer for each schedulable entity, requiring a large amount of memory for storing the state of each timer and checking each timer to see if it has expired. The state space required scales linearly with the number of schedulable entities, and can become burdensome for large numbers of schedulable entities. As the number of resource users, each corresponding to a different schedulable entity, increases, the state space required to manage the scheduler and allocate resources properly increases rapidly.

Thus, it is desirable to provide a system and method for resource scheduling capable of limiting schedulable entities to a certain minimum and maximum resource allocation, without

requiring the significant state space required by a typical rate controller module. The system should also ensure that resources are shared fairly among the different schedulable entities.

20200909 14:28:00

SUMMARY OF THE INVENTION

The present invention schedules resource requests from a plurality of schedulable entities while limiting the maximum and minimum quality of service allocated to each schedulable entity. One embodiment of a scheduler in accordance with the present invention requires less
5 memory to maintain state information than existing rate-controlling schedulers, and is thus more easily scalable to large numbers of users. The scheduler also schedules resources fairly among competing schedulable entities.

A resource scheduler uses a fair-share scheduling algorithm to select resource requests to service from multiple request queues, each associated with a schedulable entity. After a resource request is selected from a queue, a rate controller checks to ensure that servicing the request will not cause the associated user's maximum quality of service to be exceeded. If the maximum quality of service will not be exceeded, the request is serviced and the virtual time is incremented. If the maximum quality of service will be exceeded, the virtual time is still incremented, but the actual request is not serviced and remains pending.

The features and advantages described in the specification are not all-inclusive, and particularly, many additional features and advantages will be apparent to one of ordinary skill in the art in view of the drawings, specification, and claims hereof. Moreover, it should be noted that the language used in the specification has been principally selected for readability and instructional purposes, and may not have been selected to delineate or circumscribe the inventive
20 subject matter, resort to the claims being necessary to determine such inventive subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is an illustration of a resource request scheduler adapted to limit the maximum quality of service allocated to each schedulable entity.

Fig. 2 is a flowchart of the process for selecting and servicing requests from schedulable entities while limiting the maximum quality of service allocated to each schedulable entity.

Fig. 3 is an illustration of a hierarchical resource request scheduler adapted to limit the maximum quality of service allocated to each schedulable entity.

The figures depict a preferred embodiment of the present invention for purposes of illustration only. One skilled in the art will readily recognize from the following discussion that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles of the invention described herein.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Reference will now be made in detail to several embodiments of the present invention, examples of which are illustrated in the accompanying drawings. Wherever practicable, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

5 Before describing various embodiments of the present invention, some further background on scheduling methods is provided.

Any separately identifiable source of requests for a resource is referred to herein as a “schedulable entity.” As examples, a schedulable entity may represent a single individual, a group of individuals with some shared association, a computer or a set of computer programs associated with a resource. For example, a company A may contain two divisions Y and Z. If resource requests, such as requests for network bandwidth, are only identified as originating from company A, company A is a “schedulable entity,” and all resource requests from both division Y and division Z are included in the same scheduling queue. However, if company A network bandwidth requests may be separately identified and scheduled between division Y and
15 division Z, then both Y and Z are schedulable entities. Separate qualities of service may be assigned to the different divisions, and requests from each division are placed into separate scheduling queues.

Different quality of service guarantees may be assigned to different types of resources for the same schedulable entity. Many different types of resources may be assigned quality of
20 service guarantees, such as CPU time, memory access time, file access time, and networking resources. Other types of resources will be evident to one of skill in the art. For example, schedulable entity “Entity A” may have a CPU quality of service set to 50% of the physical

computer's resources. Entity A may also be assigned only 40% of the physical resource's memory access time. In another embodiment, a schedulable entity may have a single quality of service guarantee that applies to each type of resource. For example, Entity A may be assigned a minimum quality of service of 40%, and a maximum quality of service of 50% for all resources.

5 In this case, Entity A will receive between 40 and 50% of all of the resources of the physical computer.

The quality of service assigned to each schedulable entity for each resource type is stored in a quality of service table or similar data structure. The quality of service is expressed as either a minimum and maximum percentage share of resources, or as a minimum and maximum quantity of a particular resource. In one embodiment, the minimum and maximum quality of service are equal, and a single quality of service value bounds the resources allocated to a particular schedulable entity. The present invention does not limit how quality of service parameters are set or selected for entities, and so any mechanism for managing quality of service may be used.

Fair-share scheduling algorithms are well known in the art, and numerous different variations exist. Fair-share scheduling algorithms partition a resource among multiple users such that each user is allocated a fair share of the resource. For purposes of example, the start-time fair queuing algorithm with virtual time scheduling will be discussed herein. However, it will be understood by one of skill in the art that numerous other fair-share scheduling algorithms may be

20 used with the inventive techniques disclosed herein. For example, a round-robin, a deficit round-robin, or a self-clocked fair queuing algorithm may be used.

Certain fair-share scheduling algorithms use various methods to emulate the idealized scheduling discipline of generalized processor sharing (GPS). In GPS, each schedulable entity

has a separate queue. The scheduler serves each non-empty queue by servicing an infinitesimally small amount of data from each queue in turn. Each queue may be associated with a service weight, and the queue receives service in proportion to this weight. Weighted fair queuing algorithms simulate the properties of GPS by calculating the time at which a request would receive service under GPS, and then servicing requests in order of these calculated service times. The start-time fair queuing algorithm is a variation of weighted fair queuing.

A virtual time may be used in fair-share scheduling algorithms. A virtual time scheduler increments a virtual time variable. A start and finish tag is calculated for each incoming resource request, and requests are serviced in order of their tags. The virtual time is equal to the current request's start tag number, and increments as additional start tags are calculated. A virtual time scheduler simulates time-division multiplexing of requests in much the same way as weighted fair queuing algorithms simulate the properties of GPS. Virtual time scheduling, weighted fair queuing, and the start-time fair queuing algorithm are discussed in "An Engineering Approach to Computer Networking" by S. Keshav, pp. 209-263, (Addison-Wesley Professional Computing Series) (1997), and "Start-time Fair Queuing: A Scheduling Algorithm for Integrated Services Packet Switching Networks" by Pawan Goyal, Harrick M. Vin, and Haichen Cheng, IEEE/ACM Transactions on Networking, Vol. 5, No. 5, pp. 690-704 (October 1997), the subject matter of both of which are hereby incorporated in their entirety.

The start-time fair queuing algorithm is used with virtual time scheduling in the following manner. Incoming resource requests are placed in separate queues within a resource scheduler, with each queue holding the requests from a particular schedulable entity. Each resource request is tagged with both a start number (SN) and a finish number (FN) tag as it reaches the head of its respective queue, and requests are serviced in order of increasing start

numbers. The start and finish number tags are calculated for each request k from a schedulable entity i , where i represents the queue in which the request k has been placed. The start and finish number tags are calculated using the virtual time $V(t)$, the weight given to each schedulable entity $\Phi(i)$ (if weighting is being implemented) and the resource duration $D(i, k, t)$ for which each request is scheduled.

The virtual time $V(t)$ is initially zero. When the resource scheduler is busy, the virtual time at time t is defined to be equal to the start tag of the request in service at time t . At the end of a busy period, the virtual time is set to the maximum finish tag assigned to any resource request that has been serviced.

The start number tag SN of a request k arriving at an inactive queue i (one that does not currently contain previous pending requests by the queue's schedulable entity) is set to the maximum of either the current virtual time $V(t)$ or the finish number FN of the previous request $(k-1)$ pending in the queue:

$$SN(i, k, t) = \max[V(t), FN(i, k-1, t)] \quad (1)$$

The finish number tag FN of a request k is the sum of the start number SN of the request and its request duration $D(i, k, t)$ divided by its weight $\Phi(i)$:

$$FN(i, k, t) = SN(i, k, t) + \frac{D(i, k, t)}{\Phi(i)} \quad (2)$$

The weight $\Phi(i)$ for each schedulable entity corresponds to the minimum quality of service assigned to that schedulable entity. Schedulable entities are provided with resources in proportion to their weights, and thus a schedulable entity with a minimum quality of service guarantee of 40% of a resource receives proportionally more resources than another schedulable

entity with a minimum quality of service guarantee of 20% of a resource. As shown in Eq. 2, as $\Phi(i)$ increases towards 1 (100% of the resources available), the amount $D(i,k,t)/\Phi(i)$ added to the SN decreases, and thus the finish number FN tag is lower for the current request ($k-1$). As shown in Eq. 1, a lower finish number FN tag on the current request ($k-1$) causes the next request k to have a lower start number SN tag. A lower SN tag means that the next request k is serviced more quickly. Thus increasing the weight of a particular queue increases the frequency at which requests are serviced from the queue.

In one embodiment, each schedulable entity queue i has an individually assigned minimum quality of service, and thus each schedulable entity is allotted a different proportion of the overall resources. In another embodiment, all of the schedulable entities are allocated the same minimum quality of service guarantee, and weights are not implemented.

The request duration D for which a resource request is scheduled is specific to the type of resource being requested and the particular scheduling system implementation. Request duration D determines the granularity of resource scheduling. For example, assume a process P requires a duration of 100 seconds of CPU time. Scheduling process P continuously for 100 seconds would mean that all other processes would starve for 100 seconds, an undesirable result. Instead, process P will typically be scheduled for a shorter upper bound duration D_{max} , such as 20 seconds. After 20 seconds, process P is preempted and another process is scheduled. If the original duration D requested is shorter than the upper bound duration D_{max} , the entire requested duration D will be serviced. Additionally, the process may block on I/O earlier than the upper bound duration on CPU time.

Various embodiments of the present invention will next be discussed. Fig. 1 illustrates a scheduler suitable for scheduling the shared usage of a variety of different types of

resources. For example, scheduler 100 may be used to schedule resources for CPU time, memory access, disk access, or networking resources such as bus bandwidth. Additional types of resources suitable for scheduling with the scheduler 100 will be evident to one of skill in the art. Different schedulable entities make requests for resources, such as a request for CPU time, a request for memory access, a request for disk access, or a request for bandwidth to transport signals within the network.

Scheduler 100 selects resource requests for service using a fair-share scheduling algorithm. Scheduler 100 additionally prevents a selected resource request from being serviced if the maximum quality of service assigned to the schedulable entity that made the resource request would thereby be exceeded. Scheduler 100 implements rate control over the maximum quality of service in a manner that preserves the fairness properties of the scheduling algorithm. Scheduler 100 further implements rate control in a manner that reduces the memory required to store state variables as compared to prior-art methods. Scheduler 100 may be implemented as part of the operating system of a computer.

Different quality of service guarantees are implemented by allocating different amounts of the scheduled resource to servicing each of the schedulable entities. Resources may be allocated to different schedulable entities as a percentage of a particular resource, for example, allocating 50% of the resource to schedulable entity A and 25% to schedulable entity B. Resources may also be allocated as a particular number of units of a resource, for example, the operating system may be instructed to allocate x seconds of memory access to schedulable entity A and y seconds of memory access to schedulable entity B.

Each schedulable entity has an assigned minimum and maximum quality of service guarantee for the resource being scheduled. The minimum guarantee represents the minimum amount of a particular resource the schedulable entity should receive. The maximum guarantee represents the maximum amount of a particular resource the schedulable entity should receive, and this maximum will be enforced even if the resource will become idle. In one embodiment, the minimum and maximum quality of service guarantees are equal. In another embodiment, the maximum quality of service exceeds the minimum quality of service, for example, schedulable entity A is guaranteed at least 20% of the resource, but at no time will entity A receive more than 30% of the resource.

Scheduler 100 receives incoming resource requests 110 from a group of schedulable entities who share the resource being scheduled. The incoming requests 110 are sorted into separate queues, with each queue holding the pending requests for a single schedulable entity. Three queues 130, 140 and 150 are shown in Fig. 1. Queue 130 holds two pending resource requests 132A and 132B. Queue 140 holds three pending resource requests 142A, 142B, and 142C. Queue 150 holds one pending resource request 152A. It will be evident to one of skill in the art that additional queues may be added to the scheduler 100 if additional schedulable entities are to share the resource being scheduled.

Requests 132A, 142A and 152A reside in the head of queues 130, 140, and 150, respectively. As resource requests are serviced, they leave the head of the queue, and requests remaining in the queue move up in the queue. The fair-share selector 180 selects resource requests for service from the head of each queue based upon a fair-share scheduling algorithm, which ensures that each schedulable entity receives its minimum quality of service. Each resource request is assigned a start number tag SN and finish number tag FN (Eqs. 1 and 2) as it

reaches the head of its respective queue. Selector 180 selects the next resource request to service as the one with the lowest SN tag. The fair-share selector 180 will not allocate service time to an empty queue.

If the schedulable entities have equal minimum quality of service guarantees (equal weights), the fair-share scheduling algorithm apportions resources equally among the schedulable entities. If the minimum quality of service guarantees for the schedulable entities differ, each schedulable entity has a weight $\Phi(i)$ that is incorporated into the fair-share scheduling algorithm, as shown in Eq. 2. The weight $\Phi(i)$ thus influences the tags assigned to each resource request and the resulting selection order of requests.

In the following example, it will be assumed that the scheduler 100 uses the start-time fair queuing algorithm with a virtual clock tracking the virtual time $V(t)$. Selector 180 also limits each selected request to a pre-determined maximum duration D_{max} , thereby limiting the amount of resource time allocated to any single request.

As resource requests enter the head of each queue, the scheduler 100 assigns each request a start number tag SN using Eq. 1, and a finish number tag FN using Eq. 2. Selector 180 selects requests for service based upon their start number SN order, with the lowest SN selected first. Ties are broken arbitrarily. Each request includes a request duration $D_{request}$. The request will be scheduled for service for a duration $D=D_{request}$; however, if $D_{request}$ is greater than the scheduler 100's pre-determined upper bound duration D_{max} , the selector 180 will only permit D_{max} of the request to be selected for service:

$$D = \min(D_{request}, D_{max}) \quad (3)$$

If $D_{request} > D_{max}$, the remainder of the request ($D_{remainder} = D_{request} - D_{max}$) will be returned to the head of its queue, and a new start number and finish number tag will be calculated for the remainder of the request.

The virtual time $V(t)$ is related to the current request's start number SN . Each time that the selector 180 selects 114 a request for service with a new start number tag SN , this advances the virtual time $V(t)$ of the scheduler 100. As shown in Eq. 1, this calculation of SN depends on the finish number FN calculation given in Eq. 2. A component of the FN calculation is the request duration $D(i,k,t)$, and thus the request duration also influences the virtual time $V(t)$.

Once a request has been selected 114 for service, it is checked 116 by the rate controller for its respective queue. Queue 130 has a rate controller 136, queue 140 has a rate controller 146, and queue 150 has a rate controller 156. Each rate controller checks to determine whether the selected request is eligible for service, i.e. whether servicing the selected request will exceed the maximum quality of service guarantee for the associated queue's schedulable entity. Techniques for determining whether satisfying the current selected request would result in the request's schedulable entity exceeding its pre-specified maximum quality of service are well known in the art. Rate controlled schedulers are discussed in "An Engineering Approach to Computer Networking" by S. Keshav, pp. 248-252, (Addison-Wesley Professional Computing Series) (1997), the subject matter of which is herein incorporated by reference in its entirety.

If servicing the request would exceed the maximum quality of service guarantee, the request is not eligible for service and the rate controller leaves the request pending at the head of its respective queue. However, the rate controller will send a dummy request to be serviced 120 that is scheduled for a zero time duration D , thereby updating the virtual time. The SN and FN tags for the request left pending at the head of the non-eligible queue will thus be recalculated as

if the request had just arrived at the head of the queue. This SN and FN tag recalculation occurs both for requests that are not eligible for service, and for request remainders as discussed previously.

If the rate controller determines that the request is eligible for service, the request is removed from its queue. The request is then serviced, meaning that the request is allowed to consume the resource being scheduled for a duration D .

Scheduler 100 employs rate controllers 136, 146 and 156 to ensure that each schedulable entity does not exceed its maximum quality of service guarantee. The rate controllers do not require a separate set of rate controller queues, nor does each rate controller require a separate timer. Thus, the rate controllers require less memory for state variable storage as compared to prior-art rate control mechanisms.

Fig. 2 is a flowchart of the process for selecting and servicing resource requests as implemented within a resource scheduling module, such as the scheduler 100. The method of Fig. 2 will be illustrated by an example of scheduling a CPU resource using scheduler 100 of Fig. 1. For purposes of example, assume that the CPU resource's D_{max} is 20 seconds. The SN tag, and $D_{request}$ associated with each head-of-queue resource request is given in Table 1 below:

<u><i>SN tag</i></u>	<u>Resource request</u>	<u><i>D_{request}</i> (seconds)</u>
1	132A	50
3	142A	20
6	152A	30

Table 1

The scheduler reviews 200 the *SN* tags for the head-of-queue requests (132A, 142A, and 152A). The scheduler selects 210 the resource request with the smallest or “next” *SN* tag (132A). In one embodiment, ties are broken arbitrarily. In another embodiment, the system administrator specifies an order for resolving tag number ties. The scheduler then determines 215 the duration *D* to allot to the selected resource request. In this example, since $D_{request} > D_{max}$ (50 seconds > 20 seconds), the request 132A will only be granted D_{max} (20 seconds) of CPU time.

The rate controller (136) associated with the queue of the selected request (132A) is called 220. The rate controller determines 230 whether servicing the selected request will exceed the maximum quality of service guarantee allocated to the request’s schedulable entity. If the maximum quality of service will be exceeded, the scheduler sends a dummy request 250 for servicing, thereby simulating servicing the request within the fair-share scheduling algorithm controlling the request selection process. The virtual time $V(t)$ then advances to the *SN* tag of the next request, just as if the request 132A had actually been serviced. If the maximum quality of service will not be exceeded, the resource request is actually serviced 240.

The scheduler then calculates 260 new tags as needed. If the entire requested duration $D_{request}$ of the selected request (132A) was serviced 240, resource request 132B moves to the head of queue 130 and a new SN and FN tag is assigned to request 132B. However, if the entire request or part of the request (132A) was not serviced and is still pending, new tags are
5 calculated for the portion of request 132A that was not serviced. In this example, 30 seconds of original request 132A remain to be serviced. The remainder of request 132A with an updated duration $D_{request}$ of 30 seconds remains pending in the queue 130. New SN and FN tags are calculated for the remainder of request 132A, and request 132B remains back in the queue 130.

The scheduler then returns to step 200 and reviews the head-of-queue SN tags to select the next request for service.

The fair-share scheduling system and method described in Figs. 1 and 2 may also be implemented in a hierarchical fair-share scheduler. A hierarchical fair-share scheduler performs fair-share scheduling on multiple levels, i.e. schedulable entity groups may have subgroups that also have weights. A set of start number and finish number tags are maintained at each level of the hierarchy. The hierarchical fair-share scheduler is implemented as a tree structure, where each node is a scheduler that partitions the resource allocated to it and schedules child nodes. A hierarchical fair-share scheduler is discussed in "Start-time Fair Queuing: A Scheduling
13 Algorithm for Integrated Services Packet Switching Networks" by Pawan Goyal, Harrick M. Vin, and Haichen Cheng, IEEE/ACM Transactions on Networking, Vol. 5, No. 5, pp. 690-704
20 (October 1997), the subject matter of which is incorporated herein in its entirety.

A hierarchical fair-share scheduler may be used, for example, if a particular parent schedulable entity wishes to allocate a certain percentage of the parent's total resources to a first child schedulable entity, and allocate the remainder to a second schedulable entity. For example,

assume Company A requests that a minimum of 40% of the resources of a particular CPU be guaranteed for the Company A, and Company A is prohibited from using more than 45% of the resources. Company A also wishes to ensure that its Division X receives at least 60% of the Company A CPU resources, and that its Division Y receives the remainder of the resources.

5

The hierarchical resource scheduler will constrain Company A to using between 40 and 45% of the available resources of the shared CPU, using a weighted fair-share queuing algorithm and rate controllers as described in Figs. 1 and 2 to implement the desired minimum and maximum quality of service. Company A, however, may request that the resource scheduler implement one of several different methods for resource sharing between Divisions X and Y. In one embodiment, Divisions X and Y are assigned resources using a non-work-conserving scheduling algorithm wherein both X and Y are constrained to a minimum and a maximum quality of service. In another embodiment, Divisions X and Y are assigned resources using a work-conserving scheduling algorithm wherein both X and Y are guaranteed a minimum quality of service, and any additional resources are shared between X and Y according to their respective weights.

FIG. 3

20

Fig. 3 illustrates a hierarchical fair-share scheduler 300. A parent scheduler 302 has two child schedulers 301A and 301B. Parent scheduler 302 includes two parent queues 380A and 380B, corresponding to two schedulable entities. Child scheduler 301A includes two child queues 330A and 330B, corresponding to two schedulable entities, both of which feed resource requests 320A to parent queue 380A. Child scheduler 301B includes three child queues 340A, 340B and 340C, all of which feed resource requests 320B to parent queue 380B. Each parent queue represents a main schedulable entity (such as a company), and each child queue represents a subgroup schedulable entity of its parent (such as a division of the company).

Initial resource requests 310 are separated by child schedulable entity and placed into their corresponding child queues 330 or 340. Child schedulers 301A and 301B assign start and finish number tags to requests in the heads of their respective queues. Each child scheduler 301 maintains a separate set of tags, and the parent scheduler 302 also maintains a separate set of tags. Consequently, each scheduler has a separate virtual time clock. Selector 326A selects resource requests for service from requests at the head of queues 330A and 330B using a fair-share scheduling algorithm. Selector 326B also selects resource requests for service from requests at the head of queues 340A, 340B and 340C using a fair-share scheduling algorithm.

In one embodiment, each child queue is assigned a weight $\Phi(i)$ increasing or decreasing the queue's relative resource share. Each child queue also has an associated rate controller that limits the maximum resource share that that child queue may obtain. Child queue 330A has a rate controller 336A; child queue 330B has a rate controller 336B; child queue 340A has a rate controller 346A; child queue 340B has a rate controller 346B; and child queue 340C has a rate controller 346C. When a request from the head of a child queue is selected for service, the associated rate controller determines if servicing the request will exceed the maximum quality of service allocated to the child queue. If the maximum quality of service will be exceeded, a dummy request for zero resources is sent for service as a placeholder, and the request remains pending in the head of its child queue. Start and finish number tags are updated as described previously, thereby incrementing the virtual time for the scheduler.

Requests (including dummy requests) selected for service by selector 326A that are not blocked by their respective rate controllers are output 320A into queue 380A of scheduler 302. Similarly, requests selected for service by selector 326B that are not blocked by their respective rate controllers are output 320B into queue 380B of scheduler 302. Scheduler 302 assigns new

start and finish number tags to requests in queues 380A and 380B, thereby incrementing the virtual time for the parent scheduler 302. A selector 328 uses a fair-share scheduling algorithm to select requests for service from the heads of queues 380A and 380B. Queue 380A has an associated rate controller 386A, and queue 380B has an associated rate controller 386B.

5 Scheduler 302 uses the method described in Fig. 2 to output resource requests for servicing 322, subject to minimum and maximum quality of service constraints on queues 380A and 380B.

In the embodiment shown in Fig. 3, child schedulers 301A and 301B are implemented in a manner similar to the parent scheduler 302. In another embodiment, the hierarchical resource scheduler 300 implements one or more of the child schedulers 301A and 301B differently than parent scheduler 302. For example, child schedulers 301A and/or 301B may be implemented without rate controllers associated with each child queue. A child queue without a rate controller will not be subject to a maximum quality of service limitation. Additionally, child schedulers 301A and/or 301B may be implemented as work-conserving schedulers, or may use different types of scheduling algorithms.

The resource scheduling method of the present invention is suitable for use in scheduling the resources of “virtual servers.” It is desirable for an ISP to provide multiple server applications on a single physical host computer, in order to allow multiple customers to use a single host computer. If a different customer is associated with each server application, or “virtual server”, the ISP will implement a method of sharing resources between customers.

20 Additionally, it is desirable to be able to constrain virtual server customers to a minimum and maximum quality of service guarantee. This allows customers to be limited to a certain amount of the resources of the physical host computer.

The resource scheduler of the present invention may be used in the context of virtual servers to schedule some or all of the physical host computer resources. Each virtual server corresponds to a separate schedulable entity. Each virtual server is assigned a maximum and minimum quality of service guarantee. In another embodiment, separate maximum and minimum quality of service guarantees may be assigned to different resources used by the same virtual server. Requests for resources from each virtual server are serviced according to the resource scheduling method of the present invention.

Although the invention has been described in considerable detail with reference to certain embodiments, other embodiments are possible. As will be understood by those of skill in the art, the invention may be embodied in other specific forms without departing from the essential characteristics thereof. For example, different fair-share algorithms may be implemented in the resource request scheduler. Additionally, a weighted fair-share or hierarchical weighted fair-share algorithm implementation may be used in the resource request scheduler. Accordingly, the present invention is intended to embrace all such alternatives, modifications and variations as fall within the spirit and scope of the appended claims and equivalents.

We claim:

1 1. A method for scheduling a resource to service a plurality of pending requests received
2 from a plurality of schedulable entities, while preventing each schedulable entity from exceeding
3 a maximum quality of service allocated to each schedulable entity, comprising:

4 selecting a request associated with a schedulable entity;

5 responsive to determining that servicing the selected request will exceed the
6 schedulable entity's maximum quality of service, advancing a virtual time for
 scheduling the requests, without servicing the request; and

 responsive to determining that servicing the selected request does not exceed the
 schedulable entity's maximum quality of service, servicing the request and
 advancing the virtual time.

2. The method of claim 1, wherein the request includes a request to allocate disk space.

3. The method of claim 1, wherein the request includes a request to allocate memory.

1 4. The method of claim 1, wherein the request includes a request for network bandwidth.

1 5. The method of claim 1, wherein the request includes a request for CPU processing
2 cycles.

1 6. The method of claim 1, wherein the request is selected using a fair-share scheduling
2 algorithm.

1 7. The method of claim 6, wherein the fair-share scheduling algorithm is a weighted fair-
2 share scheduling algorithm, each weight corresponding to a schedulable entity's minimum
3 quality of service allocation.

1 8. The method of claim 7, wherein the minimum quality of service allocated to each
2 schedulable entity is a minimum percentage share of the resource.

1 9. The method of claim 6, wherein the fair-share scheduling algorithm is a hierarchical
2 fair-share scheduling algorithm.

1 10. The method of claim 6, wherein the fair-share scheduling algorithm is a hierarchical
2 weighted fair-share scheduling algorithm, each weight corresponding to a schedulable entity's
3 minimum quality of service allocation.

1 11. The method of claim 6, wherein the fair-share scheduling algorithm is a start-time
2 fair queuing algorithm with virtual time scheduling.

1 12. The method of claim 1, wherein each request includes a requested duration, the
2 method further including:

3 limiting the requested duration of the request to a pre-determined request duration
4 upper bound.

1 13. The method of claim 1, wherein the maximum quality of service allocated to each
2 schedulable entity is a maximum percentage share of the resource.

1 14. The method of claim 1, wherein a rate controller determines if servicing the request
2 will exceed the schedulable entity's maximum quality of service.

1 15. The method of claim 14, wherein if the rate controller determines that servicing the
2 request will exceed the schedulable entity's maximum quality of service, the request remains
3 pending.

1 16. A method for scheduling resource requests from a plurality of schedulable entities,
2 wherein each resource request includes a requested duration and each schedulable entity has a
3 maximum quality of service guarantee, the method comprising:

4 assigning a start number tag to a resource request using a start-time fair queuing
5 algorithm with virtual time scheduling;

6 selecting the resource request with the smallest start number tag, the selected request
7 having an associated schedulable entity;

8 limiting the requested duration of the selected resource request to a pre-determined
9 duration upper bound;

10 servicing the selected resource request if servicing the selected resource request will
11 not exceed the associated schedulable entity's maximum quality of service
12 guarantee; and

13 advancing a virtual time value.

1 17. The method of claim 16, further including:

2 updating the start number tag for a resource request associated with the schedulable
3 entity that made the selected resource request if the selected resource request
4 is not serviced.

1 18. The method of claim 16, further including:

2 leaving the selected resource request pending if servicing the selected resource
3 request will exceed the schedulable entity's maximum quality of service
4 guarantee.

1 19. A system for scheduling pending resource requests from a plurality of schedulable
2 entities while limiting a maximum quality of service allocated to each schedulable entity,
3 comprising:

4 a plurality of schedulable entity queues for holding pending resource requests, each
5 schedulable entity queue holding resource requests from a schedulable entity;

6 a scheduler for selecting resource requests from the plurality of schedulable entity queues
7 using a fair-share scheduling algorithm, and further adapted to increment a virtual time value
8 each time a resource request is selected; and

9 a plurality of rate controllers associated with the plurality of schedulable entity queues,
10 each rate controller adapted to limit the rate at which resource requests selected by the scheduler
11 are serviced to the schedulable entity's maximum quality of service.

1 20. The system of claim 19, wherein each rate controller is further adapted to:

2 monitor the servicing of resource requests from the rate controller's associated
3 schedulable entity queue to calculate the quality of service received by the
4 schedulable entity; and

5 block the servicing of a selected resource request if the schedulable entity's maximum
6 quality of service would be exceeded if the selected resource request was
7 serviced.

8 21. The system of claim 19, wherein each schedulable entity queue is associated with a
9 weight, and the scheduler uses a weighted fair-share queuing algorithm.

10 22. A hierarchical system for scheduling resource requests from a plurality of child
11 schedulable entities while limiting the maximum quality of service allocated to a plurality of
12 parent schedulable entities, comprising:

13 a plurality of child schedulable entity queues for holding pending resource requests, each
14 child schedulable entity queue holding resource requests from a child schedulable entity;

15 one or more child schedulers for selecting resource requests from the plurality of child
16 schedulable entity queues using a fair-share scheduling algorithm, and further adapted to
17 transmit selected resource requests to a parent schedulable entity queue;

18 a plurality of parent schedulable entity queues, each parent schedulable entity queue
19 receiving resource requests from a subset of the child schedulable entity queues, each parent
20 schedulable entity queue holding resource requests received from one of the child schedulers;

12 a parent scheduler for selecting resource requests from the plurality of parent schedulable
13 entity queues using a fair-share scheduling algorithm, and further adapted to increment a virtual
14 time value each time a resource request is selected; and

15 a plurality of rate controllers associated with the plurality of parent schedulable entity
16 queues, each rate controller adapted to limit the rate at which resource requests selected by the
17 parent scheduler are serviced to a parent schedulable entity's maximum quality of service.

1 23. A computer program product for scheduling a plurality of pending requests for
2 service from a resource received from a plurality of schedulable entities, while preventing each
3 schedulable entity from exceeding a maximum quality of service allocated to each schedulable
4 entity, the computer program product comprising:

5 a computer readable medium that stores program code including:

6 program code that selects a request associated with a schedulable entity using a
7 fair-share scheduling algorithm;

8 program code that services the request if a rate controller determines that
9 servicing the request will not exceed the associated schedulable entity's
10 maximum quality of service; and

11 program code that advances a virtual time in the fair-share scheduling algorithm.

1 24. The computer program product of claim 23, wherein the fair-share scheduling
2 algorithm is a weighted fair-share scheduling algorithm, each weight corresponding to a
3 schedulable entity's minimum quality of service allocation.

1 25. The computer program product of claim 23, wherein each request includes a
2 requested duration, the computer program product further including:

3 program code that limits the requested duration of the request to a pre-determined
4 request duration upper bound.

1

FAIRLY PARTITIONING RESOURCES WHILE LIMITING THE MAXIMUM FAIR SHARE

ABSTRACT OF THE DISCLOSURE

Resource requests from a plurality of schedulable entities are scheduled while limiting
5 the maximum and minimum quality of service allocated to each schedulable entity. The resource
scheduler of the present invention requires less memory to maintain state information than
existing rate-controlling schedulers, and is thus more easily scalable to large numbers of users.
The resource scheduler also schedules resources fairly among competing schedulable entities. A
fair-share scheduling algorithm is used by a resource scheduler to select resource requests to
10 service. A rate controller checks to ensure that servicing the selected request will not cause the
associated user's maximum quality of service to be exceeded. If the maximum quality of service
will not be exceeded, the virtual time used in the scheduling algorithm is incremented, and the
request is serviced. If the maximum quality of service will be exceeded, the virtual time is still
incremented, but the request is not serviced and remains pending.

Scheduler 100




Figure 1

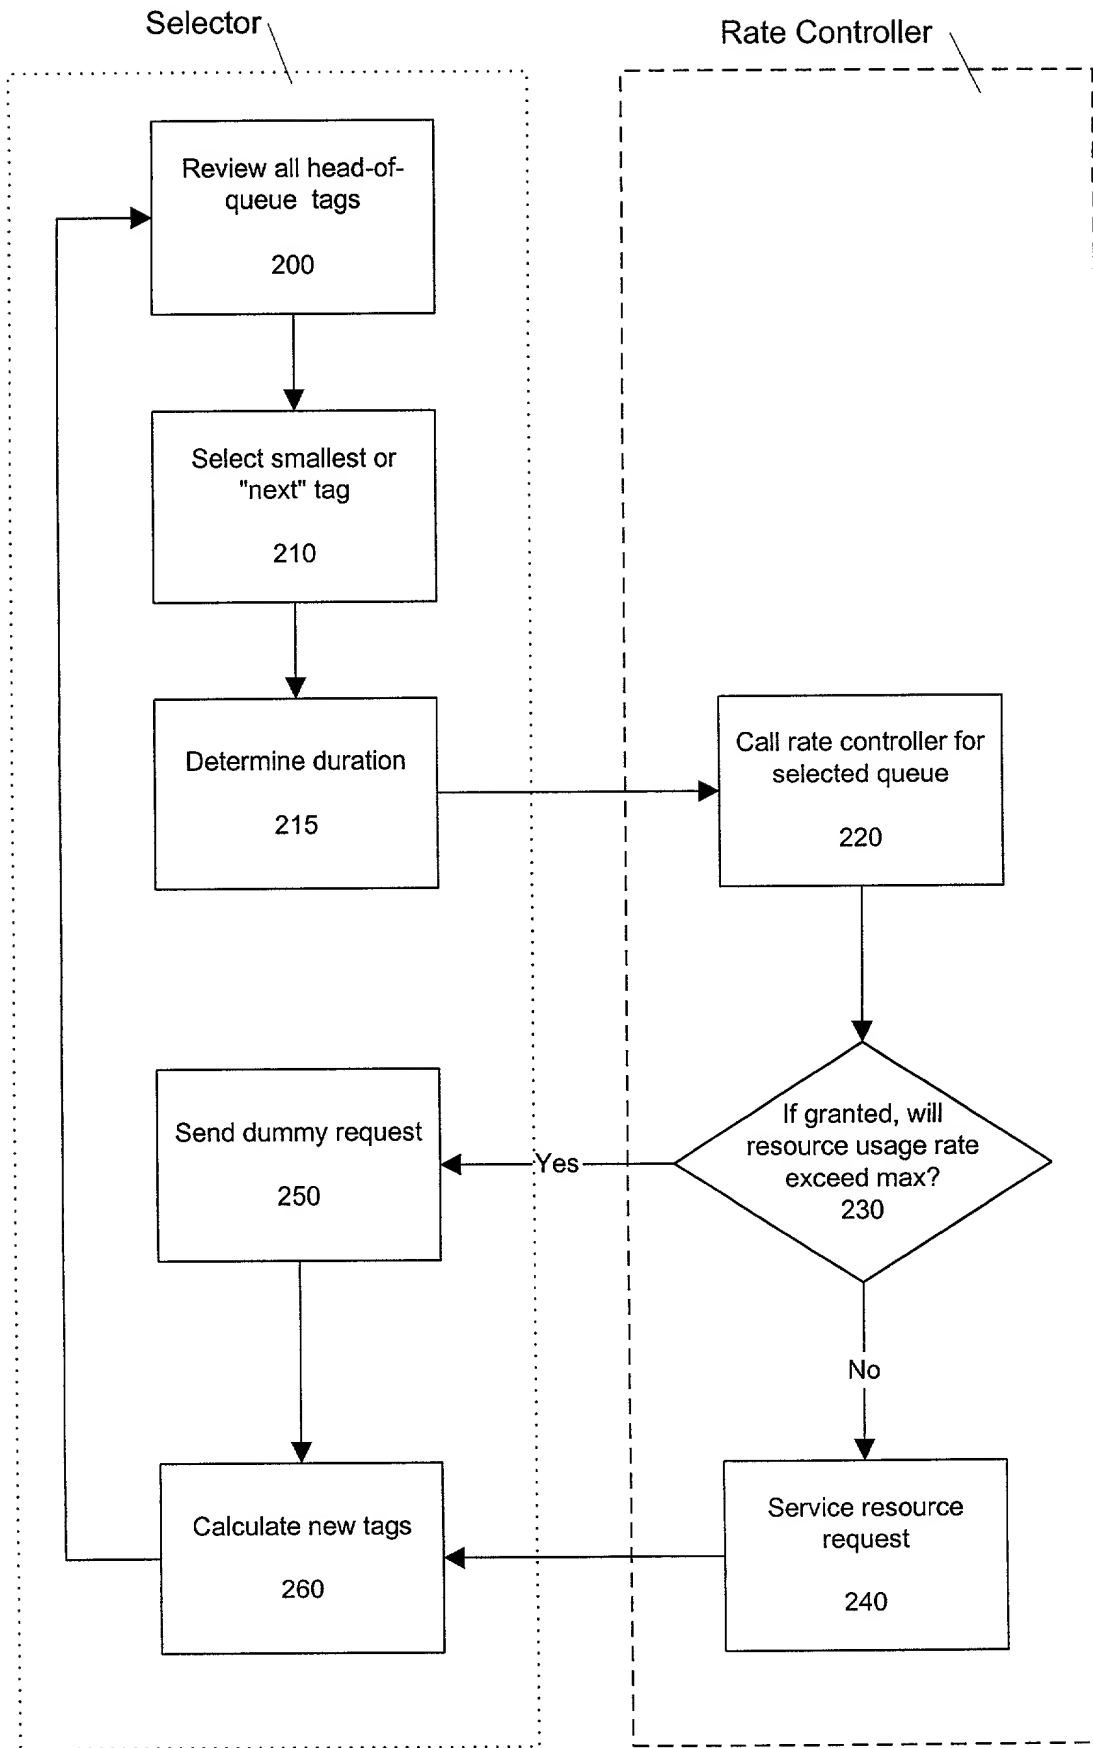


Figure 2

2025 RELEASE UNDER E.O. 14176




Figure 3

0010/PTO Rev. 6/95	U.S. Department of Commerce Patent and Trademark Office	Attorney Docket Number	4466 US
DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION		First Named Inventor	Pawan Goyal
		COMPLETE IF KNOWN	
		Application Number	not yet known
		Filing Date	not yet known
		Group Art Unit	not yet known
		Examiner Name	not yet known
<input checked="" type="checkbox"/> Declaration Submitted with Initial Filing OR <input type="checkbox"/> Declaration Submitted after Initial Filing			

As a below named inventor, I hereby declare that:

My residence, post office address, and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

Fairly Partitioning Resources While Limiting The Maximum Fair Share

the specification of which

(Title of the Invention)

☒ is attached hereto

OR

☐ was filed on (MM/DD/YYYY) [] as United States Application Number or PCT International Application Number [] and was amended on (MM/DD/YYYY) [] (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in Title 37 Code of Federal Regulations. § 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code § 119 (a)-(d) or § 385(b) of any foreign application(s) for patent or inventor's certificate, or § 365 (a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? YES NO	
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

☐ Additional foreign application numbers are listed on a supplemental priority sheet attached hereto:

I hereby claim the benefit under Title 35, United States Code § 119(e) of any United States provisional application(s) listed below.

Application Number(s)	Filing Date (MM/DD/YYYY)	<input type="checkbox"/> Additional provisional application numbers are listed on a supplemental sheet attached hereto.

DECLARATION	Page 2
--------------------	---------------

I hereby claim the benefit under Title 35, United States Code § 120 of any United States application(s), or § 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of Title 35, United States Code § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. Parent Application Number	PCT Parent Number	Parent Filing Date (MM/DD/YYYY)	Parent Patent Number (if applicable)

☐ Additional U.S. or PCT international application numbers are listed on a supplemental priority sheet attached hereto.

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

Name	Registration Number	Name	Registration Number
Albert C. Smith Laura A. Majerus	20,355 33,417	Robert R. Sachs Renée M. DuBord	42,120 42,500

☐ Additional attorney(s) and/or agent(s) named on a supplemental sheet attached hereto.

Please direct all correspondence to:

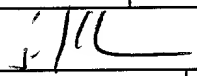
Robert R. Sachs
Fenwick & West LLP
Two Palo Alto Square
Palo Alto, CA 94306
U.S.A.

Telephone (650) 858-7110	Fax (650) 494-1417
--------------------------	--------------------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Name of Sole or First Inventor:		<input type="checkbox"/> A petition has been filed for this unsigned inventor					
Given Name	Pawan	Middle Initial		Family Name	Goyal	Suffix e.g. Jr.	
Inventor's Signature					Date	8/2/00	
Residence: City	Mountain View	State	CA	Country	U.S.A.	Citizenship	India
Mailing Address	777 W. Middlefield Road, #83						
Mailing Address							
City	Mountain View	State	CA	Zip	94043	Country	U.S.A.

☒ Additional inventors are being named on supplemental sheet(s) attached hereto

DECLARATION				ADDITIONAL INVENTOR(S) Supplemental Sheet			
Name of Additional Joint Inventor, if any:				<input checked="" type="checkbox"/> A petition has been filed for this unsigned inventor			
Given Name	Srinivasan	Middle Initial		Family Name	Keshav	Suffix e.g. Jr.	
Inventor's Signature					Date	8/2/00	
Residence: City	Mountain View	State	CA	Country	U.S.A.	Citizenship	India
Mailing Address							
Mailing Address	834 Sutter Avenue						
City	Mountain View	State	CA	Zip	94043	Country	U.S.A.

Name of Additional Joint Inventor, if any:				<input checked="" type="checkbox"/> A petition has been filed for this unsigned inventor			
Given Name		Middle Initial		Family Name		Suffix e.g. Jr.	
Inventor's Signature					Date		
Residence: City		State		Country		Citizenship	
Mailing Address							
Mailing Address							
City		State		Zip		Country	.

Name of Additional Joint Inventor, if any:				<input checked="" type="checkbox"/> A petition has been filed for this unsigned inventor			
Given Name		Middle Initial		Family Name		Suffix e.g. Jr.	
Inventor's Signature					Date		
Residence: City		State		Country		Citizenship	
Mailing Address							
Mailing Address							
City		State		Zip		Country	

Name of Additional Joint Inventor, if any:				<input checked="" type="checkbox"/> A petition has been filed for this unsigned inventor			
Given Name		Middle Initial		Family Name		Suffix e.g. Jr.	
Inventor's Signature					Date		
Residence: City		State		Country		Citizenship	
Mailing Address							
Mailing Address							
City		State		Zip		Country	

☐ Additional inventors are being named on supplemental sheet(s) attached hereto